
A Statistical Test for Joint Distributions Equivalence

Francesco Solera
University of Modena, Italy
francesco.solera@unimore.it

Andrea Palazzi
University of Modena, Italy
andrea.palazzi@unimore.it

Abstract

We provide a distribution-free test that can be used to determine whether any two joint distributions p and q are statistically different by inspection of a large enough set of samples. Following recent efforts from Long *et al.* [1], we rely on joint kernel distribution embedding to extend the *kernel two-sample test* of Gretton *et al.* [2] to the case of joint probability distributions. Our main result can be directly applied to verify if a dataset-shift has occurred between training and test distributions in a learning framework, without further assuming the shift has occurred only in the input, in the target or in the conditional distribution.

1 Introduction

Detecting when dataset shifts occur is a fundamental problem in learning, as one needs to re-train its system to adapt to the new data before making wrong predictions. Strictly speaking, if $(X, Y) \sim p$ is the training data and $(X', Y') \sim q$ is the test data, a dataset shift occurs when the hypothesis that (X, Y) and (X', Y') are sampled from the same distribution wanes, that is $p \neq q$. The aim of this work is to provide a statistical test to determine whether such a shift has occurred given a set of samples from training and testing set.

To cope with the complexity of joint distribution, a lot of literature has emerged in recent years trying to approach easier versions of the problem, where the distributions were assumed to differ only by a factor. For example a covariate shift is met when, in the decomposition $p(x, y) = p(y|x)p(x)$, $p(y|x) = q(y|x)$ but $p(x) \neq q(x)$. Prior distribution shift, conditional shift and others can be defined in a similar way. For a good reference, the reader may want to consider Quiñero-Candela *et al.* [3] or Moreno-Torres *et al.* [4].

2 Preliminaries

As it often happens, such assumptions are too strong to hold in practice and do require an expertise about the data distribution at hand which cannot be given for granted. A recent work by Long *et al.* [1] has tried to tackle the same question, but without making restricting hypothesis on what was changing between training and test distributions. They developed the *Joint Distribution Discrepancy* (JDD), a way of measuring distance between any two joint distributions – regardless of everything else. They build on the *Maximum Mean Discrepancy* (MMD) introduced in Gretton *et al.* [2] by noticing that a joint distribution can be mapped into a tensor product feature space via kernel embedding.

The main idea behind MMD and JDD is to measure distance between distributions by comparing their embeddings in a Reproducing Kernel Hilbert Space (RKHS). RKHS \mathcal{H} is a Hilbert space of functions $f : \Omega \mapsto \mathbb{R}$ equipped with inner products $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ and norms $\| \cdot \|_{\mathcal{H}}$. In the context of this work, all elements $f \in \mathcal{H}$ of the space are probability distributions that can be evaluated by means of inner products $f(x) = \langle f, k(x, \cdot) \rangle_{\mathcal{H}}$ with $x \in \Omega$, thanks to the reproducing property. k is a kernel function that takes care of the embedding by defining an implicit feature mapping

$k(x, \cdot) = \phi(x)$, where $\phi : \Omega \mapsto \mathcal{H}$. As always, $k(x, x') = \langle \phi(x), \phi(x') \rangle_{\mathcal{H}}$ can be viewed as a measure of similarity between points $x, x' \in \Omega$. If a characteristic kernel is used, then the embedding is injective and can uniquely preserve all the information about a distribution [5]. According to the seminal work by Smola *et al.* [6], the kernel embedding of a distribution $p(x)$ in \mathcal{H} is given by $\mathbb{E}_x[k(x, \cdot)] = \mathbb{E}_x[\phi(x)] = \int_{\Omega} \phi(x) dP(x)$.

Having all the required tools in place, we can introduce the MMD and the JDD.

Definition 1 (Maximum Mean Discrepancy (MMD) [2]). *Let $\mathcal{F} \subset \mathcal{H}$ be the unit ball in a RKHS. If x and x' are samples from distributions p and q respectively, then the MMD is*

$$\text{MMD}(\mathcal{F}, p, q) = \sup_{f \in \mathcal{F}} (\mathbb{E}_x[f(x)] - \mathbb{E}_{x'}[f(x')]) \quad (1)$$

Definition 2 (Joint Distribution Discrepancy (JDD) [1]). *Let $\mathcal{F} \subset \mathcal{H}$ be the unit ball in a RKHS. If (x, y) and (x', y') are samples from joint distributions p and q respectively, then the JDD is*

$$\begin{aligned} \text{JDD}(\mathcal{F}, p, q) &= \sup_{f, g \in \mathcal{F}} (\mathbb{E}_{x, y}[f(x)g(y)] - \mathbb{E}_{x', y'}[f(x')g(y')]) \\ &= \|\mathbb{E}_{x, y}[\phi(x) \otimes \psi(y)] - \mathbb{E}_{x', y'}[\phi(x') \otimes \psi(y')]\|_{\mathcal{F} \otimes \mathcal{F}}, \end{aligned} \quad (2)$$

where ϕ and ψ are the mappings yielding to kernels k_{ϕ} and k_{ψ} , respectively.

Note that, conversely to Long *et al.* [1], we don't square the norm in Eq. (2). A biased empirical estimation of JDD can be obtained by replacing the population expectation with the empirical expectation computed on samples $\{(x_1, y_1), (x_2, y_2), \dots, (x_m, y_m)\} \in X \times Y$ from p and samples $\{(x'_1, y'_1), (x'_2, y'_2), \dots, (x'_n, y'_n)\} \in X' \times Y'$ from q :

$$\begin{aligned} \text{JDD}_b(\mathcal{F}, X, Y, X', Y') &= \sup_{f, g \in \mathcal{F}} \left(\frac{1}{m} \sum_{i=1}^m f(x_i)g(y_i) - \frac{1}{n} \sum_{i=1}^n f(x'_i)g(y'_i) \right) \\ &= \left\| \frac{1}{m} \sum_{i=1}^m \phi(x_i) \otimes \psi(y_i) - \frac{1}{n} \sum_{i=1}^n \phi(x'_i) \otimes \psi(y'_i) \right\|_{\mathcal{F} \otimes \mathcal{F}} \\ &= \left(\frac{1}{m^2} \sum_{i, j=1}^m k_{\phi}(x_i, x_j) k_{\psi}(y_i, y_j) + \frac{1}{n^2} \sum_{i, j=1}^n k_{\phi}(x'_i, x'_j) k_{\psi}(y'_i, y'_j) + \dots \right. \\ &\quad \left. - \frac{2}{mn} \sum_{i=1}^m \sum_{j=1}^n k_{\phi}(x_i, x'_j) k_{\psi}(y_i, y'_j) \right)^{\frac{1}{2}} \end{aligned} \quad (3)$$

Moreover, throughout the paper, we restrict ourselves to the case of bounded kernels, specifically $0 \leq k(x_i, x_j) \leq K$, for all i and j and for all kernels.

3 The test

Under the null hypothesis that $p = q$, we would expect the JDD to be zero and the empirical JDD to be converging towards zero as more samples are acquired. The following theorem provides a bound on deviations of the empirical JDD from the ideal value of zero. These deviations may happen in practice, but if they are too large we will want to reject the null hypothesis.

Theorem 1. *Let p, q, X, X', Y, Y' be defined as in Sec. 1 and Sec. 2. If the null hypothesis $p = q$ holds, and for simplicity $m = n$, we have*

$$\text{JDD}_b(\mathcal{F}, X, X', Y, Y') \leq \sqrt{\frac{8K^2}{m} (2 - \log(1 - \alpha))} \quad (4)$$

with probability at least α . As a consequence, the null hypothesis $p = q$ can be rejected with a significance level α if Eq. 4 is not satisfied.

Interestingly, Type II errors probability decreases to zero at rate $\mathcal{O}(m^{\frac{1}{2}})$ – preserving the same convergence properties found in the *kernel two-sample test* of Gretton *et al.* [2]. We warn the reader that this result was obtained by neglecting dependency between X and Y . See Sec. 4 and following for a deeper discussion.

4 Experiments

To validate our proposal, we handcraft joint distribution starting from MNIST data as follows. We sample an image i from a specific class and define the pair of observation (x_i, y_i) as the vertical and horizontal projection histograms of the sampled image. Fig. 1(a) depicts the process. The number of samples obtained in the described manner is defined by m and they all belong to the same class. It is easy to see why the distribution is joint. For all the experiments we employed an RBF kernel, which is known to be characteristic [7], *i.e.* induces a one-to-one embedding. Formally, for (x, y) and (x', y') distributed according to p or q indistinctly, we have

$$k_\phi(x, x') = \exp\left(-\frac{\|x - x'\|^2}{\sigma_\phi^2}\right) \quad \text{and} \quad k_\psi(y, y') = \exp\left(-\frac{\|y - y'\|^2}{\sigma_\psi^2}\right). \quad (5)$$

The parameters σ_ϕ and σ_ψ have been experimentally set to 0.25. Accordingly, both kernels are bounded by $K = 1$.

In the first experiment we obtain $(X, Y) = \{(x_i, y_i)\}_i$ by sampling $m = 1000$ images from the class of number 3. Similarly, we collect $(X', Y') = \{(x'_i, y'_i)\}_i$ by applying a rotation ρ to other m sampled images from the same class. Of course, when $\rho = 0$, the two samples (X, Y) and (X', Y') come from the same distribution and the null hypothesis that $p = q$ should not be rejected. On the opposite, as ρ increases in absolute value we expect to see JDD increase as well – up to the point of exceeding the critical value defined in Eq. (4). Fig. 1(b) illustrates the behavior of JDD when (X, Y) and (X', Y') are sampled from increasingly different distributions.

To deepen the analysis, in Fig. 2 we study the behavior of the critical value by changing the significance level α and the sample size m .




Figure 1: In (a) we show an exemplar image drawn from the MNIST dataset. Observation (x, y) are the projection histograms along both axis, *i.e.* x (y) is obtained by summing values across rows (columns). On the right, (b) depicts the behavior of the JDD measure when samples (X', Y') are drawn from a different distribution w.r.t. (X, Y) , specifically the distribution of rotated images. The rotation is controlled by the ρ parameter. The green line shows the critical value for rejecting the null hypothesis (acceptance region below).

5 Limitations and conclusions

The proof of Theorem 1 is based on the McDiarmid’s inequality which is not defined for joint distributions. As a result, we considered all random variables of both distributions independent of each others, despite being clearly rarely the case. However, empirical (but preliminary) experiments show encouraging results, suggesting that the test could be safely applied to evaluate the equivalence of joint distributions under broad independence cases.

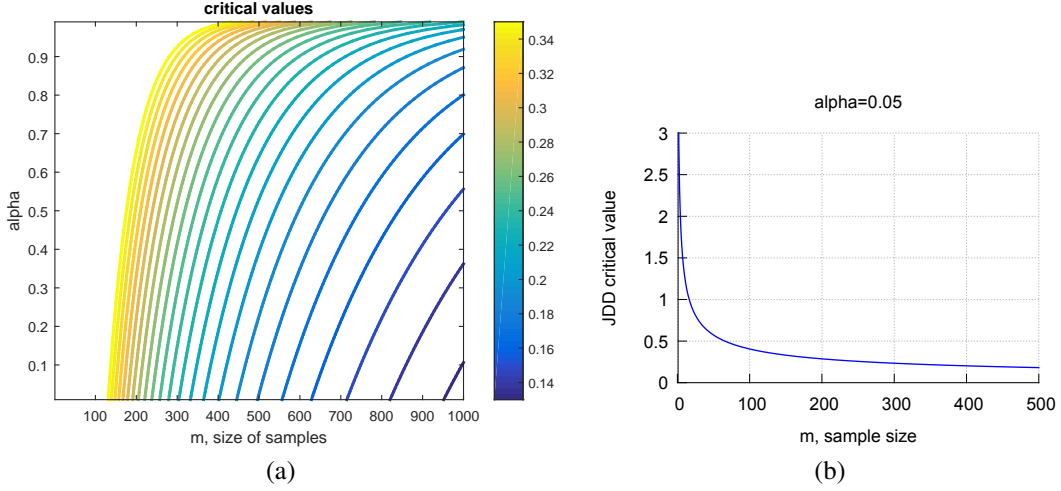


Figure 2: On the left, (a) depicts the JDD critical value for the test of Eq. (4) when the significance level α and the sample size m change. Cooler colors correspond to lower values of the JDD threshold. Not surprisingly, more conclusive and desirable tests can be obtained either by lowering α or by increasing m . Complementary, (b) shows the convergence rate of the test threshold at increasing size m of sample, for a fix value of $\alpha = 0.05$. It is worth noticing, that the elbow of the convergence curve is found around $m = 50$.

6 Appendices

6.1 Preliminaries to the proofs

In order to prove our test, we first need to introduce McDiarmid's inequality and a modified version of Rademacher average with respect to the m -sample (X, Y) obtained from a joint distribution.

Theorem 2 (McDiarmid's inequality [8]). *Let $f : \mathcal{X}^m \rightarrow \mathbb{R}$ be a function such that for all $i \in \mathbb{N}_m$, there exist $c_i < \infty$ for which*

$$\sup_{X \in \mathcal{X}^m, \tilde{x} \in \mathcal{X}} |f(x_1, \dots, x_m) - f(x_1, \dots, x_{i-1}, \tilde{x}, x_{i+1}, \dots, x_m)| \leq c_i. \quad (6)$$

Then for all probability measures p and every $\xi > 0$,

$$\Pr_X(f(X) - \mathbf{E}_X[f(X)] > \xi) < \exp\left(-\frac{2\xi^2}{\sum_{i=1}^m c_i^2}\right), \quad (7)$$

where \mathbf{E}_X denotes the expectation over the m random variables $x_i \sim p$, and \Pr_X denotes the probability over these m variables.

Definition 3 (Joint Rademacher average). *Let \mathcal{F} be the unit ball in an RKHS on the domain $\mathcal{X} \times \mathcal{Y}$, with kernels bound between 0 and K . Let $(X, Y) = \{(x_1, y_1), (x_2, y_2), \dots, (x_m, y_m)\}$ be an i.i.d. sample drawn according to probability measure p on $\mathcal{X} \times \mathcal{Y}$, and let σ_i be i.i.d. and taking values in $\{-1, +1\}$ with equal probability. We define the joint Rademacher average*

$$\hat{R}_m(\mathcal{F}, X, Y) = \mathbf{E}_\sigma \left[\sup_{f, g \in \mathcal{F}} \left| \frac{1}{m} \sum_{i=1}^m \sigma_i f(x_i) g(y_i) \right| \right]. \quad (8)$$

Theorem 3 (Bound on joint Rademacher average). *Let $\hat{R}_m(\mathcal{F}, X, Y)$ be the joint Rademacher average defined as in Def. 3, then*

$$\hat{R}_m(\mathcal{F}, X, Y) \leq K/m^{\frac{1}{2}}. \quad (9)$$

Proof. The proof follows the main steps from Bartlett and Mendelson [9], lemma 22. Recall that $f(x_i) = \langle f, \phi(x_i) \rangle$ and $g(y_i) = \langle f, \psi(y_i) \rangle$, for all x_i and y_i .

$$\begin{aligned}
\hat{R}_m(\mathcal{F}, X, Y) &= \mathbf{E}_\sigma \left[\sup_{f, g \in \mathcal{F}} \left| \frac{1}{m} \sum_{i=1}^m \sigma_i f(x_i) g(y_i) \right| \right] \\
&= \frac{1}{m} \mathbf{E}_\sigma \left[\left\| \sum_{i=1}^m \sigma_i \phi(x_i) \otimes \psi(y_i) \right\| \right] \\
&\leq \frac{1}{m} \left(\sum_{i,j=1}^m \mathbf{E}_\sigma [\sigma_i \sigma_j k_\phi(x_i, x_j) k_\psi(y_i, y_j)] \right)^{\frac{1}{2}} \\
&= \frac{1}{m} \left(\sum_i^m \mathbf{E}_\sigma [\sigma_i^2 k_\phi(x_i, x_i) k_\psi(y_i, y_i)] \right)^{\frac{1}{2}} \\
&= \frac{1}{m} \left(\sum_i^m k_\phi(x_i, x_i) k_\psi(y_i, y_i) \right)^{\frac{1}{2}} \\
&\leq \frac{1}{m} (mK^2)^{\frac{1}{2}} = K/m^{\frac{1}{2}}
\end{aligned} \tag{10}$$

□

6.2 Proof of Theorem 1

We start by applying McDiarmid's inequality to JDD_b under the simplifying hypothesis that $m = n$,

$$\text{JDD}_b = \sup_{f, g \in \mathcal{F}} \left(\frac{1}{m} \sum_{i=1}^m f(x_i) g(y_i) - f(x'_i) g(y'_i) \right).$$

Without loss of generality, let us consider the variation of JDD_b with respect to any x_i . Since \mathcal{F} is the unit ball in the Reproducing Kernel Hilbert Space we have

$$|f(x_i)| = |\langle f, \phi(x_i) \rangle| \leq \|f\| \|\phi(x_i)\| \leq 1 \times \sqrt{\langle \phi(x_i), \phi(x_i) \rangle} = \sqrt{k(x_i, x_i)} \leq \sqrt{K} \tag{11}$$

for all $f \in \mathcal{F}$ and for all x_i . Consequently, the largest variation to JDD_b is bounded by $2K/m$, as the bound in Eq. 11 also holds for all $g \in \mathcal{F}$ and for all y_i . Summing up squared maximum variations for all x_i, y_i, x'_i and y'_i , the denominator in Eq. (7) becomes

$$4m \left(\frac{2K}{m} \right)^2 = \frac{16K^2}{m}, \tag{12}$$

yielding to

$$\Pr_{X,Y,X',Y'}(\text{JDD}_b - \mathbf{E}_{X,Y,X',Y'}[\text{JDD}_b] > \xi) < \exp \left(-\frac{m\xi^2}{8K^2} \right). \tag{13}$$

To fully exploit McDiarmid's inequality, we also need to bound the expectation of JDD_b . To this end, similarly to Gretton *et al.* [2], we exploit symmetrisation (Eq. (14)(d)) by means of a *ghost sample*, i.e. a set of observations whose sampling bias is removed through expectation (Eq. (14)(b)).

In particular, let (\bar{X}, \bar{Y}) and (\bar{X}', \bar{Y}') be i.i.d. samples of size m drawn independently of (X, Y) and (X', Y') respectively, then

$$\begin{aligned}
\mathbf{E}_{X,Y,X',Y'}[\text{JDD}_b] &= \mathbf{E}_{X,Y,X',Y'} \left[\sup_{f,g \in \mathcal{F}} \left(\frac{1}{m} \sum_{i=1}^m f(x_i)g(y_i) - \frac{1}{m} \sum_{i=1}^m f(x'_i)g(y'_i) \right) \right] \\
&\stackrel{(a)}{=} \mathbf{E}_{X,Y,X',Y'} \left[\sup_{f,g \in \mathcal{F}} \left(\frac{1}{m} \sum_{i=1}^m f(x_i)g(y_i) - \mathbf{E}_{x,y}[fg] - \frac{1}{m} \sum_{i=1}^m f(x'_i)g(y'_i) + \mathbf{E}_{x',y'}[fg] \right) \right] \\
&\stackrel{(b)}{=} \mathbf{E}_{X,Y,X',Y'} \left[\sup_{f,g \in \mathcal{F}} \left(\frac{1}{m} \sum_{i=1}^m f(x_i)g(y_i) - \mathbf{E}_{\bar{X},\bar{Y}} \left[\frac{1}{m} \sum_{i=1}^m f(\bar{x}_i)g(\bar{y}_i) \right] - \dots \right. \right. \\
&\quad \left. \left. \frac{1}{m} \sum_{i=1}^m f(x'_i)g(y'_i) + \mathbf{E}_{\bar{X}',\bar{Y}'} \left[\frac{1}{m} \sum_{i=1}^m f(\bar{x}'_i)g(\bar{y}'_i) \right] \right) \right] \\
&\stackrel{(c)}{\leq} \mathbf{E}_{X,Y,X',Y',\bar{X},\bar{Y},\bar{X}',\bar{Y}'} \left[\sup_{f,g \in \mathcal{F}} \left(\frac{1}{m} \sum_{i=1}^m f(x_i)g(y_i) - \frac{1}{m} \sum_{i=1}^m f(\bar{x}_i)g(\bar{y}_i) - \dots \right. \right. \\
&\quad \left. \left. \frac{1}{m} \sum_{i=1}^m f(x'_i)g(y'_i) + \frac{1}{m} \sum_{i=1}^m f(\bar{x}'_i)g(\bar{y}'_i) \right) \right] \\
&\stackrel{(d)}{=} \mathbf{E}_{X,Y,X',Y',\bar{X},\bar{Y},\bar{X}',\bar{Y}',\sigma,\sigma'} \left[\sup_{f,g \in \mathcal{F}} \left(\frac{1}{m} \sum_{i=1}^m \sigma_i(f(x_i)g(y_i) - f(\bar{x}_i)g(\bar{y}_i)) + \dots \right. \right. \\
&\quad \left. \left. \frac{1}{m} \sum_{i=1}^m \sigma'_i(f(x'_i)g(y'_i) - f(\bar{x}'_i)g(\bar{y}'_i)) \right) \right] \\
&\leq \mathbf{E}_{X,Y,\bar{X},\bar{Y},\sigma} \left[\sup_{f,g \in \mathcal{F}} \left| \frac{1}{m} \sum_{i=1}^m \sigma_i(f(x_i)g(y_i) - f(\bar{x}_i)g(\bar{y}_i)) \right| \right] + \dots \\
&\quad \mathbf{E}_{X',Y',\bar{X}',\bar{Y}',\sigma'} \left[\sup_{f,g \in \mathcal{F}} \left| \frac{1}{m} \sum_{i=1}^m \sigma'_i(f(x'_i)g(y'_i) - f(\bar{x}'_i)g(\bar{y}'_i)) \right| \right] \\
&\leq 2\mathbf{E}_{X,Y}[\hat{R}_m(\mathcal{F}, X, Y)] + 2\mathbf{E}_{X',Y'}[\hat{R}_m(\mathcal{F}, X', Y')] \\
&\leq 4K/m^{\frac{1}{2}}.
\end{aligned} \tag{14}$$

In Eq. (14), (b) adds a difference $\mathbf{E}_{x,y}[fg] - \mathbf{E}_{x',y'}[fg]$ that equals 0 since $p = q$ by the null hypothesis, and (c) employs Jensen's inequality.

By substituting the upper bound of Eq. (14) in Eq. (13), we obtain Theorem 1.

References

- [1] Mingsheng Long, Jianmin Wang, and Michael I Jordan. Deep transfer learning with joint adaptation networks. *arXiv preprint arXiv:1605.06636*, 2016.
- [2] Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13(Mar):723–773, 2012.
- [3] Joaquin Quionero-Candela, Masashi Sugiyama, Anton Schwaighofer, and Neil D Lawrence. *Dataset shift in machine learning*. The MIT Press, 2009.
- [4] Jose G Moreno-Torres, Troy Raeder, Rocío Alaiz-Rodríguez, Nitesh V Chawla, and Francisco Herrera. A unifying view on dataset shift in classification. *Pattern Recognition*, 45(1):521–530, 2012.
- [5] Kenji Fukumizu, Arthur Gretton, Xiaohai Sun, and Bernhard Schölkopf. Kernel measures of conditional dependence. In *NIPS*, volume 20, pages 489–496, 2007.

- [6] Alex Smola, Arthur Gretton, Le Song, and Bernhard Schölkopf. A hilbert space embedding for distributions. In *International Conference on Algorithmic Learning Theory*, pages 13–31. Springer, 2007.
- [7] Arthur Gretton, Karsten M Borgwardt, Malte Rasch, Bernhard Schölkopf, and Alex J Smola. A kernel method for the two-sample-problem. In *Advances in neural information processing systems*, pages 513–520, 2006.
- [8] Colin McDiarmid. On the method of bounded differences. *Surveys in combinatorics*, 141(1):148–188, 1989.
- [9] Peter L Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002.